

AP168HO

**APPLICATION  
FOR  
UNITED STATES LETTERS PATENT**

Be it known that we, William Chen, of 138 Rock Harbor Lane, Foster City, CA 94404, a citizen of the United States of America and Jau-Yuen Chen, of 420 James Road, Apt. 17, Palo Alto, CA 94306, a citizen of Taiwan - Republic of China, have invented new and useful improvements in:

**METHOD AND APPARATUS FOR SUMMARIZING AND INDEXING  
THE CONTENTS OF AN AUDIO-VISUAL PRESENTATION**

of which the following is the specification

CERTIFICATION UNDER 37 C.F.R. 1.10

"Express Mail" Mailing Label Number: EV311301795US

Date of Deposit: September 26, 2003

I hereby certify that this patent application is being deposited with the United States Postal Service on this date in an envelope as "Express Mail Post Office to Addressee" service under 37 C.F.R. 1.10 on the date indicated above and is addressed to Mail Stop Patent Application, Commissioner for Patents, P.O. Box 1450, Alexandria, VA 22313-1450.

  
\_\_\_\_\_  
Ann F. George

# METHOD AND APPARATUS FOR SUMMARIZING AND INDEXING THE CONTENTS OF AN AUDIO- VISUAL PRESENTATION

*by Inventors*

William Chen and Jau-Yuen Chen

## BACKGROUND OF THE INVENTION

### 1. Field of the Invention

10 [0001] This invention relates generally to information processing and more particularly to a method and apparatus for summarizing and indexing the contents of an audiovisual presentation.

### 2. Description of the Related Art

[0002] Formal presentations serve an important and popular means of communication. In academia and industry, the capture of such presentations for subsequent online viewing has become routine for applications such as distance-learning and technical training. Recording the seminars and placing the content online provides users the benefits of anywhere, anytime, and anyone viewing due to the ubiquitous nature of the Internet. Additionally, seminars having multiple presentations running concomitantly, force an individual to make a choice to attend one of the multiple presentations, when the individual may desire to attend more than one of the concomitant presentations.

[0003] Previous work on automatic video summarization may be characterized to fall into one of three broad areas: segmentation, analysis, and presentation. Segmentation involves the partitioning of a frame of video. For the domain of audio-visual presentations, segmentation requires consideration of changing lighting conditions, speaker movements, and camera pan/zooms. One of the shortcomings of the available

techniques that segment audio-visual presentations is the inability to effectively handle the changing lighting conditions, speaker movements, and camera pan/zooms. Furthermore, there are no available techniques that are capable of indexing the audio visual content once the content is segmented, nor are there any techniques for  
5 summarizing the content for easy retrieval by a user. The problems become more acute when a user is accessing the video data through a handheld device with limited computational resources.

[0004] As a result, there is a need to solve the problems of the prior art to enable automatic indexing and an effective scheme for summarizing the content of an  
10 audiovisual presentation that enables a user to efficiently locate desired information.

**SUMMARY OF THE INVENTION**

[0005] Broadly speaking, the present invention fills these needs by providing a method and system capable of automatically summarizing the contents of an audiovisual presentation in real time. It should be appreciated that the present invention can be implemented in numerous ways, including as a method, a system, computer readable media or a device. Several inventive embodiments of the present invention are described below.

[0006] In one embodiment, a method for segmenting image data is provided. The method initiates with identifying a pixel associated with a current frame of the image data. Then, a neighborhood of pixels is defined around the pixel associated with the current frame. The defining of the neighborhood includes generating a three dimensional neighborhood. Next, a distance between the pixel associated with the current frame and each pixel associated with the neighborhood of pixels is compared to determine a smallest distance. Then, if the pixel associated with the current frame belongs to a current segment of the image data is determined based upon the smallest distance.

[0007] In another embodiment, a method for creating a summary of an audiovisual presentation is provided. The method initiates with segmenting a frame of the audiovisual presentation. Then, a slide region of the segmented frame is identified. Next, a histogram representing lines in the slide region is generated. Then, moving regions associated with successive frames from the histogram are suppressed.

[0008] In yet another embodiment, a computer readable media having program instructions for segmenting image data is provided. The computer readable medium includes program instructions for identifying a pixel associated with a current frame of the image data. Program instructions for defining a neighborhood of pixels around the

pixel associated with the current frame are provided, where the program instructions for defining the neighborhood include program instructions for program instructions for generating a three dimensional neighborhood. Program instructions for comparing a distance between the pixel associated with the current frame and each pixel associated  
5 with the neighborhood of pixels to determine a smallest distance are included. Program instructions for determining if the pixel associated with the current frame belongs to a current segment of the image data based upon the smallest distance are also included.

[0009] In still yet another embodiment, a computer readable medium having program instructions for creating a summary of an audiovisual presentation is provided. The  
10 computer readable medium includes program instructions for segmenting a frame of the audiovisual presentation. Program instructions for identifying a slide region of the segmented frame are provided. Program instructions for generating a histogram representing lines in the slide region and program instructions for suppressing moving regions associated with successive frames from the histogram are included.

15 [0010] In another embodiment, a system configured to capture and summarize an audiovisual presentation is provided. The system includes a recording device capable of capturing audio and video signals from the presentation. A computing device in communication with the recording device is included. The computing device has access to audiovisual data of the audiovisual presentation. The computing device includes a  
20 slide segmentation module configured to extract a slide region from a frame of the video signals according to a single pass color segmentation scheme.

[0011] In yet another embodiment, a system configured to provide a real-time summarization of a meeting is provided. The system includes an image capture device configured to capture a presentation associated with the meeting. A media server  
25 configured to receive captured presentation data from the image capture device is

included. The media server has access to copies of presentation media used for the meeting. The media server is further configured to generate summary data corresponding to the presentation from the captured presentation data. The summary data is associated with presentation media transition points of the meeting. A client in communication with  
5 the media server is also included. The client is capable of receiving the summary data.

[0012] In still yet another embodiment, an integrated circuit is provided. The integrated circuit includes segmentation circuitry configured to segment a frame of image data into regions. The segmentation circuitry is capable of identifying one of the regions as a slide region through analysis of a color characteristic and a shape characteristic associated with  
10 each of the regions. Shot detection circuitry configured to identify a group of frames associated with the frame through analysis of edge information of the slide region with adjacent frames of the image data is also included.

[0013] Other aspects and advantages of the invention will become apparent from the following detailed description, taken in conjunction with the accompanying drawings,  
15 illustrating by way of example the principles of the invention.

**BRIEF DESCRIPTION OF THE DRAWINGS**

[0014] The present invention will be readily understood by the following detailed description in conjunction with the accompanying drawings, and like reference numerals  
5 designate like structural elements.

[0015] Figure 1 is a high-level block diagram illustrating the modules associated with the generation of the table of contents for an audio-visual presentation in accordance with one embodiment of the invention.

[0016] Figure 2 is a schematic diagram illustrating how a traditional image segmentation  
10 system is restricted to comparing pixels with four predecessors in causal order.

[0017] Figure 3 is a schematic diagram representing a technique for comparing a reference pixel with five neighbors from a current frame and a previous frame in causal order in accordance with one embodiment of the invention.

[0018] Figure 4 is an exemplary representation of a scan line order when processing a  
15 frame of video data in accordance with one embodiment of the invention.

[0019] Figures 5A through 5C represent the segmentation results from the color segmentation scan described with reference to Figure 3 and Table 1.

[0020] Figure 6 is a schematic diagram illustrating the modules for generating a one-bit representation for a slide region in accordance with one embodiment of the invention.

20 [0021] Figure 7 is a more detailed schematic diagram of the motion suppression module of Figure 6 in accordance with one embodiment of the invention.

[0022] Figure 8 represents a pictorial illustration of the motion mask in accordance with one embodiment of the invention.

[0023] Figure 9 is a video trace representing slide transitions during various frames of the  
25 video presentation in accordance with one embodiment of the invention.

[0024] Figure 10 is a schematic diagram representing a template matching module in accordance with one embodiment of the invention.

[0025] Figure 11 is a high level schematic diagram of a system capable of capturing and summarizing video from a presentation and emailing the summary to a user.

- 5 [0026] Figure 12 is a flow chart representing the method operations for creating a summary of an audio visual presentation in accordance with one embodiment of the invention.



**DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS**

[0027] An invention is described for a system, and method for automatically generating a summarization of an audiovisual presentation. It will be apparent, however, to one skilled in the art, in light of this disclosure, that the present invention may be practiced without some or all of these specific details. In other instances, well known process operations have not been described in detail in order not to unnecessarily obscure the present invention.

[0028] The embodiments described herein provide a method and system that captures and automatically summarizes an audio-visual presentation in real-time. From the video, audio, and slide presentation, a table of contents (TOC) that highlights key topics with links to the corresponding slides and video files is automatically generated. Thus, with access to the audio-visual recording of a presentation, i.e., slide presentation, and the stored presentation material, a TOC is built for the presentation so that a user may select a particular segment of the presentation. Additionally, the summarization of the presentation through the TOC, or some other suitable summarization technique, enables a user having a handheld device, e.g., a personal digital assistant (PDA), a cellular phone, a web tablet, etc., to view the summarization page. Thereafter, the user may download a specific frame of video which can be processed by the limited resources of the handheld device, as opposed to a video stream of the presentation, which would be beyond the limited computational capabilities of consumer handheld devices.

[0029] As will be explained in more detail below, the key modules in the system includes a slide segmentation module, a shot detection module, and a template matching module. The slide segmentation module is configured to extract a slide region from each frame of the digital recording device, e.g., camcorder, recording the presentation. The shot detection module then identifies groups of video frames according to slide transitions.

The template matching module then links the stored slide to a corresponding video shot through analysis of the extracted slide region.

[0030] The input into the system is a combination of audio-visual signals generated from a recording e.g., digital camcorder or any other suitable digital video recorder, of the presentation or meeting and textual information from an original presentation media, such as a slide presentation. In one embodiment, the slide presentation is a POWERPOINT presentation. For the input to the system, it is assumed that access to the audio-visual recording and the original presentation media from the presentation is available. Additionally, the projected slides of the presentation media are captured by the digital video recording.

[0031] Figure 1 is a high-level block diagram illustrating the modules associated with the generation of the table of contents for an audio-visual presentation in accordance with one embodiment of the invention. A frame of video 100 is received by slide segmentation module 102. Slide segmentation module 102 is configured to extract a slide from the frame of video 100 for template matching as will be explained in more detail below. As can be seen, slide segmentation module 102 is associated with module 110 which locates the slide region. In one embodiment, slide segmentation is performed by first applying color segmentation to each frame of the video. The slide region is then identified as the dominant, coherent color region with a compact shape (e.g., a rectangular shape ratio for a slide).

[0032] Shot detection module 104 of Figure 1, then compares successive frames of video data for differences in order to identify all frames of a segment of the video data that are associated with the slide extracted from slide segmentation module 102. As will be explained further below, shot detection module 104 is associated with module 112 where slide transitions are detected in order to identify the segment of the video data having a

same slide for each frame. In essence, shot detection module 104 parses the video into shots based on slide transitions. Each shot effectively captures the speaker presenting the contents from a single slide. In one embodiment, for robustness, the slide region is transformed into a one-bit representation using edge detection and binary thresholding.

5 The one-bit representation is then transformed to the Hough parameter domain and an edge histogram is generated from the Hough parameters. Correlation between edge histograms is used to generate a trace of the slide similarity. Peaks in this trace are used to detect slide transitions and shot boundaries. Included in shot detection module 104 is a motion suppression module configured to reduce the effects of moving objects, e.g., the  
10 speaker or an object controlled by the speaker, that intersects the slide region and causes false slide transitions.

[0033] Still referring to Figure 1, a key frame which represents the segment of the video data captured by shot detection module 104 is then matched with stored slide 108 through template matching module 106. That is, a matching algorithm which is linked to the  
15 original slides through module 114, processes the key frame data from shot detection module 104 in order to determine and match with one of original slides 108. Here, a keyframe, which contains just the extracted slide region, is used as a template and matched against each of the original slides (or copies of the original slides). For robustness, the matching algorithm preprocesses the keyframe and original slides into a  
20 one-bit representation using edge detection and binary thresholding. The one-bit representation is then transformed into an edge histogram using spatial X/Y projection. Similarity between the keyframe and the original slides is measured using a chi-squared metric on the corresponding edge histograms. Thereafter, a summary of the video presentation is generated through module 116.

[0034] Figure 2 is a schematic diagram illustrating how a traditional image segmentation system is restricted to comparing pixels with four predecessors in causal order. Here, pixel (i,j) 120e is associated with a frame of video data that includes predecessor labels associated with pixel locations 120a through 120d. However, the labels associated with pixel locations 120f through 120i are unknown at this time. Therefore, the traditional method may use a two-step process where the first step scans the frame to get the data and then a second scan is used to do the segmentation. Alternatively, the traditional process may ignore the future data of pixel locations 120f through 120i. Under either alternative, causal constraints restrict the traditional image segmentation system to only comparing pixels with the four predecessors in causal order as described above.

[0035] Figure 3 is a schematic diagram representing a technique for comparing a reference pixel with five neighbors from a current frame and a previous frame in causal order in accordance with one embodiment of the invention. Here, a three-dimensional neighborhood is created where two dimensions are represented in the current frame, i.e., the x and y coordinates, and one dimension is represented in time, i.e., the previous frame. Thus, spatial and temporal characteristics are considered in the neighborhood. It should be appreciated that the pixels associated with locations 122a and 122b are from a current frame as they are known. Whereas, the pixels associated with positions 122c, 122d, and 122e are from a previous frame. Thus, the pixels associated with positions 122c through 122e borrow future information from a previous frame. One skilled in the art will appreciate that the previous frame pixels do not differ significantly as compared to the corresponding current frame pixels, therefore, the previous frame pixels act as a good approximation for the color segmentation technique described herein.

[0036] Figure 4 is an exemplary representation of a scan line order when processing a frame of video data in accordance with one embodiment of the invention. Here, the scan

line initiates in the upper left pixel of frame 123 and zigzags through the whole frame as represented in Figure 4. It will be apparent to one skilled in the art that the scan line order is shown for exemplary purposes only and is not meant to be limiting as any suitable scan line order may be utilized. Additionally, a neighborhood of five predecessors is exemplary and not meant to be limiting as any suitable number of predecessors may be utilized. Additionally, a neighborhood of five predecessors is exemplary and not meant to be limiting as any suitable number of predecessors may be utilized. Additionally, a neighborhood of five predecessors may be utilized with the embodiments herein.

[0037] Table 1 illustrates a one pass segmentation algorithm configured to utilize the five predecessors mentioned above for video segmentation.

TABLE 1

```

10      Symbol:label(i,j;k) : the label for pixel (i,j) in frame k
      Initial label(i,j;0) to 0 for all i,j
      For each frame k in video,
      Compute centroid of each segment
      Reset number of point for each segment
15      Begin,
          For pixel (i,j) in frame k,
              Compute distance from label (i,j;k-1) to label (i-1,j;k) and label (i,j-
              l;k)
              Merge labels if distance <th3
20              Compute distance from pixel(i,j;k) to label of its causal
              predecessor as
              {label(i,j-l;k), label(i-1,j;k), label(i,j;k-1), label (i,j+1;k-1),
              label(i+1,j;k-1) }
              Let mind be smallest distance and minl be the corresponding label
25              If mind<th1,
                  Label(i,j;k)=minl;
              Else
                  Let mind2 be the smallest distance of pixel(i,j;k) and all
                  labels
30                  Let minl2 be such a label.
                  If mind2<th2;
                      Let min1=minl2
                  Else
                      Create new segment
35                      Minl = label of new segment
                      Endif
              Endif
              Update segment indexed by minl to include pixel(i,j;k)
      End
40      End.

```

Starting from the upper left pixel, the one pass algorithm zigzags through the whole frame as described in Figure 4. For each pixel  $(i,j;k)$ , the algorithm compares the distance between pixel  $(i,j;k)$ , and a label of the pixel's causal predecessor as described in Figure 3. It should be appreciated that the phrase "distance between pixels," as used  
5 herein, refers to the Euclidean distance between corresponding pixels. Of course, the color model associated with the pixels is taken into account for the distance calculation. In one embodiment, the smallest distance is then compared with a threshold to decide if pixel  $(i,j;k)$  belongs to the current segment. If the pixel doesn't belong to the current segment, the algorithm will check through all label  $i$  and compare with a lower threshold  
10 to decide if the pixel belongs to a previous segment. In effect, this reduces the number of small, isolated segments that are created because the segments are not connected. If the pixel still doesn't belong to any segment a new segment which contains only pixel  $(i,j;k)$  is then created.

[0038] A merge process is commonly used after the first round of segmentation in image  
15 segmentation. However, the merge process is usually designed in an iterative approach. To avoid the unpredictable time that the traditional merge process requires, the embodiments described herein employs a merge-in-time approach, which merges segments in the next frame by comparing  $\text{label}(i,j,k-1)$  from the previous frame with  $\text{label}(i-1,j,k)$  and  $\text{label}(i,j-1,k)$  from the current frame. This approach safely merges  
20 fragments without the risk of a long execution time.

[0039] Figures 5A through 5C represent the segmentation results from the color segmentation scan described with reference to Figure 3 and Table 1. Figure 5A represents a frame of video data 124. The frame of video data 124 includes a slide presentation. The slide presentation may include artifacts, such as reflection 128 from a  
25 projector. Figure 5B represents the frame of video data 124 from Figure 5A after which

the color segmentation technique, i.e., slide segmentation, described above has been applied. The regions having different shading within Figure 5B are identified through the color segmentation technique. For example, region 126a represents one dominant and coherent region, while region 126b represents another dominant and coherent region.

5 Region 126c represents yet another dominant and coherent region. Figure 5C represents the extracted slide region 126b from Figure 5B. That is, through the application of a one pass segmentation algorithm, such as the algorithm of TABLE 1, the slide region may be identified. As mentioned above, each of the segmented regions may be identified as a dominant coherent color type. The slide region may then be identified from the

10 remaining region by using a shape ratio configured to identify the slide region. Additionally, a threshold value may be used to discard small regions prior to checking for the shape ratio. It will be apparent to one skilled in the art that the extracted slide region 126b has poor contrast due to the dark lighting conditions and contains artifacts such as the reflection 128 from a projector.

15 **[0040]** Figure 6 is a schematic diagram illustrating the modules for generating a one-bit representation for a slide region in accordance with one embodiment of the invention. Rather than using the extracted slide region for direct comparison to a database slide in order to find a match, the extracted slide region is cleaned up through the modules of Figure 6 in order to more efficiently match the extracted slide region with a stored

20 presentation slide. The contrast associated with slide region 126b is stretched in module 132. In one embodiment, a luminance histogram is generated over the slide region and stretched at the two endpoints of the histogram until it covers the range from 0-255. It will be apparent to one skilled in the art that the contrast stretch of module 132 greatly increases the sharpness of the slide content. An edge detector is then applied to the

25 generated luminance histogram in module 134. In one embodiment, the edge detector is

a canny edge detector, however, the edge detector may be any suitable edge detector. One skilled in the art will appreciate that the edge detector of module 134 is configured to capture the important outlines of the text and figures in the slide region. The output of the edge detector is then the input to module 136 where the edges and lines of the one-bit representation is transformed into the parameter space using the Hough transform.

[0041] As is generally known, the Hough transform is a popular method of extracting geometric primitives. With respect to the embodiments described herein, there is an interest in the outlines of text and figures within the slide region. The Hough transform converts the lines from x-y spatial domain into the  $(\rho, \theta)$  parameter domain according to the following equation:

$$\rho = x \cdot \cos(\theta) + y \cdot \sin(\theta) \quad (1)$$

[0042] Here,  $\rho$  is the distance from the line to the origin, and  $\theta$  is the angle between the axis x and the perpendicular to the line vector that points from the origin to the line. Because every pixel in the image may belong to several lines, an accumulator  $A(\rho, \theta)$  that measures the strength of line parameters  $(\rho, \theta)$  is maintained. The accumulator values are then thresholded to distinguish between lines and noise features. Then, a one-dimensional histogram is generated from the accumulator to represent the lines in the slide region.

[0043] It should be appreciated that in addition to the poor contrast and lighting in slide region 126b of Figure 5C, a speaker may be moving in front of the slide region, thereby occluding text regions and creating shadows. The occlusion and shadows may create edges and lines that are captured by the Hough transform. Therefore, in order to compensate for the occlusion and shadows a motion mask is developed through motion suppression module 138. The motion mask detects moving regions and then suppresses them from the edge histogram 140 as described below.



[0044] Figure 7 is a more detailed schematic diagram of the motion suppression module of Figure 6 in accordance with one embodiment of the invention. Slide region 126b is delivered to module 150 where the frame difference of the luminance channels between adjacent frames is determined. Additionally, the binary thresholding of the frame difference is computed and the results are used to generate a silhouette. The output of module 150 is delivered to module 152 where the silhouette is copied into a separate image and assigned the value of a most recent timestamp. A time delta is set such that pixels that fall below the threshold are set to zero. It should be appreciated that this composite motion history image (MHI) now contains regions of motions grouped together by their time stamps. The composite motion history image for module 152 is then delivered to module 154 where a downward stepping flood fill is used to group and segment the most recent motion regions into motion mask 156. It will be apparent to one skilled in the art that edges located within the motion mask are now excluded from the edge histogram with reference to Figure 6 through the motion suppression module.

[0045] Figure 8 represents a pictorial illustration of the motion mask in accordance with one embodiment of the invention. Here, successive frames of video 142, 144, and 146 include slide region 126b where a presenter's hand is moving over slide region 126b. As can be seen, hand image 144a through 144n moves in a downward direction through successive frames of the video data, thereby occluding portions of slide region 126b. The motion suppression modules, with reference to Figure 7, are used to generate motion mask 156 of Figure 8. Thus, the hand movement through the successive frames is captured enabling the embodiments described herein to disregard the artifacts introduced through the motion of the hand image. In one embodiment, the motion suppression module 138 assists in suppressing false slide transitions as will be explained further below.

[0046] Figure 9 is a video trace representing slide transitions during various frames of the video presentation in accordance with one embodiment of the invention. Here, peaks 160a through 160g, and corresponding video frames 160a-1 through 160g-1, illustrate transition points where a slide is being changed. Thus, the slide regions associated with video frames 160a-1 through 160g-1 represent key frames which may be used as a template to link an original slide to the corresponding video shot. The edge histograms from adjacent frames are compared using a correlation measure as described in Equation

(2):

$$Corr = \frac{\sum_i q_i v_i}{\sqrt{\sum_i q_i^2 \sum_i v_i^2}}$$

The correlation values derived from equation 2 are used to generate a video trace, and peaks in the trace correspond to shot transitions. It should be appreciated that motion suppression helps to reduce the false peaks between frames 3000-4000 in Figure 9 by eliminating the moving regions from the correlation comparison.

[0047] Figure 10 is a schematic diagram representing a template matching module in accordance with one embodiment of the invention. Here, slide region 126b is processed through histogram stretching module 162, which functions similarly to contrast stretching module 132 with reference to Figure 6. The output of module 162 is delivered to module 134 where edge detection is performed as described above. The output of edge detection module 134 is then delivered to spatial projection module 164. Here, a one-dimensional histogram is generated by projecting edge magnitudes onto the x and y axis. In order to compare the histogram against those of the original presentation media, e.g., slides, similar processing is performed on images generated from the slides. That is, the edge detection, spatial comparison and correlation comparison are performed with the original

presentation media. Then, the original slide is matched with the slide region, i.e., key frame, that most closely correlates through the correlation comparison. It should be appreciated that the original presentation media does not need to have histogram stretching applied as the original presentation media or a copy thereof is of a sufficient  
5 quality level.

[0048] Figure 11 is a high level schematic diagram of a system capable of capturing and summarizing video from a presentation and emailing the summary to a client or a user. Image capture device 173 captures a video image of presentation 170. The captured video data is transmitted to laptop computer 172. Laptop computer 172 may be  
10 configured to execute the slide segmentation, shot detection, and template matching modules described above. Of course, laptop computer 172 may be any suitable computing device that is configured to execute the functionality described herein. Laptop computer 172 is in communication with media server 174. In one embodiment, laptop computer segments the video into shots that correspond to the original slides of the  
15 presentation. The video shots are then encoded, e.g., into a Motion Picture Expert Group (MPEG) or some other suitable audio video compression standard, and stored on media server 174. In another embodiment, a web page summary structured as table of contents 178 is created and stored on media server 174.

[0049] Still referring to Figure 11, table of contents 178 includes a number of indices  
20 where each index includes title of the slide 178a, thumbnail of the slide 178c and key frame 178b that links to the corresponding video stream. Thus, the stored web page may be emailed from media server 174 to a user having a computing device (client) configured to receive the emailed data. For example, personal digital assistant (PDA) 176, laptop 180, or any other suitable device capable of receiving email may be the  
25 recipient of the web page. Once the client receives the web page, the user can quickly

browse the TOC to get an overview of the presentation. The user may also access a full screen version of thumbnail 178c through the thumbnail or download the corresponding video shot through key frame 178b. It will be apparent to one skilled in the art that a receiving device with limited resources, i.e., a handheld electronic device, can now view  
5 the key frame or video shot, as opposed to receiving the entire video stream, which is likely beyond the capabilities of the handheld device. In one embodiment, the automated summarization technique described herein may be performed at media server 174 rather than laptop 172.

[0050] It should be appreciated that the above described embodiments may be  
10 implemented in software or hardware. One skilled in the art will appreciate that the modules may be embodied as a semiconductor chip that includes logic gates configured to provide the functionality discussed above. For example, a hardware description language (HDL), e.g., VERILOG, can be employed to synthesize the firmware and the layout of the logic gates for providing the necessary functionality described herein to  
15 provide a hardware implementation of the automatic summarization techniques and associated functionality.

[0051] Figure 12 is a flow chart representing the method operations for creating a summary of an audio visual presentation in accordance with one embodiment of the invention. The method initiates with operation 190 where a frame of the audio visual  
20 presentation is segmented. Here, the color segmentation technique described above with reference to Figures 3-5C may be used to segment the frame of the audiovisual presentation into dominant and coherent regions. The method then advances to operation 192 where a slide region of the segmented frame is identified. Here, certain characteristics, i.e., such as shape ratio, are be used to identify the slide region.

Additionally, a threshold value may be used in order to eliminate small regions of the video frame.

[0052] The method of Figure 12 then proceeds to operation 194 where a histogram representing lines in the slide region is generated. Here, the shot detection module may be used to generate the histogram. The method then moves to operation 196 where moving regions associated with successive frame are suppressed from the histogram. In one embodiment, motion suppression is applied as described above to reduce the effects of moving objects intersecting the slide region and creating false alarms during shot detection. Additionally, template matching may be performed on the histogram in order to match the slide region with a stored original or copy of the slide through correlation comparison. Thus, the video frame containing the slide region and correlated original or copy of the slide are used to create a summarization, such as the summarization with reference to Figure 11.

[0053] In summary, the above described invention provides a real time summarization of an audiovisual presentation. The summarization enables users to browse a lengthy seminar or presentation and view specific content quickly and efficiently. Additionally, the recorded content may be stored on a server, thereby enabling a user access through the Internet. The summarization enables clients with limited resources the ability to view certain shots of the presentation, where the client would be unable to otherwise process the full video stream. Thus, a video recording device may be used to capture the presentation and transmit the captured data to a computer having access to the slides used for the presentation. Through the slide segmentation module, shot detection module, and the template matching module, a summarization of the presentation is provided. In one embodiment, the summarization is in the form of a table of contents.

[0054] With the above embodiments in mind, it should be understood that the invention may employ various computer-implemented operations involving data stored in computer systems. These operations include operations requiring physical manipulation of physical quantities. Usually, though not necessarily, these quantities take the form of electrical or magnetic signals capable of being stored, transferred, combined, compared, and otherwise manipulated. Further, the manipulations performed are often referred to in terms, such as producing, identifying, determining, or comparing.

[0055] The above described invention may be practiced with other computer system configurations including hand-held devices, microprocessor systems, microprocessor-based or programmable consumer electronics, minicomputers, mainframe computers and the like. The invention may also be practiced in distributing computing environments where tasks are performed by remote processing devices that are linked through a communications network.

[0056] The invention can also be embodied as computer readable code on a computer readable medium. The computer readable medium is any data storage device that can store data which can be thereafter read by a computer system. The computer readable medium also includes an electromagnetic carrier wave in which the computer code is embodied. Examples of the computer readable medium include hard drives, network attached storage (NAS), read-only memory, random-access memory, CD-ROMs, CD-Rs, CD-RWs, magnetic tapes, and other optical and non-optical data storage devices. The computer readable medium can also be distributed over a network coupled computer system so that the computer readable code is stored and executed in a distributed fashion.

[0057] Although the foregoing invention has been described in some detail for purposes of clarity of understanding, it will be apparent that certain changes and modifications may be practiced within the scope of the appended claims. Accordingly, the present

AP168HO

embodiments are to be considered as illustrative and not restrictive, and the invention is not to be limited to the details given herein, but may be modified within the scope and equivalents of the appended claims. In the claims, elements and/or steps do not imply any particular order of operation, unless explicitly stated in the claims.

5 *What is claimed is:*